

Quasi light fields: extending the light field to coherent radiation

Anthony Accardi and Gregory Wornell

Massachusetts Institute of Technology

July 30, 2009

Imaging technologies such as dynamic viewpoint generation are engineered for incoherent radiation using the traditional light field, and for coherent radiation using electromagnetic field theory. We present a model of coherent image formation that strikes a balance between the utility of the light field and the comprehensive predictive power of Maxwell's equations. We synthesize research in optics and signal processing to formulate, capture, and form images from *quasi light fields*, which extend the light field from incoherent to coherent radiation. Our coherent cameras generalize the classic beamforming algorithm in sensor array processing, and invite further research on alternative notions of image formation.

1 Introduction

The light field represents radiance as a function of position and direction, thereby decomposing optical power flow along rays. The light field is an important tool used in many imaging applications in different disciplines, but is traditionally limited to incoherent light. In computer graphics, a rendering pipeline can compute new views at arbitrary camera positions from the light field [1]. In computational photography, a camera can measure the light field and later generate images focused at different depths, after the picture is taken [2]. In electronic displays, an array of projectors can present multiple viewpoints encoded in the light field, enabling 3D television [3]. Many recent incoherent imaging innovations have been made possible by expressing image pixel values as appropriate integrals over light field rays.

For coherent imaging applications, the value of decomposing power by position and direction has long been recognized without the aid of a light field, since the complex-valued scalar field encodes direction in its phase. A hologram encodes multiple viewpoints, but in a different way than the light field [4]. An ultrasound machine generates images focused at different depths, but from air pressure instead of light field measurements [5]. A Wigner distribution function models the operation of optical systems in simple ways, by conveniently

inferring direction from the scalar field instead of computing non-negative light field values [6]. Comparing these applications, coherent imaging uses the scalar field to achieve results similar to those that incoherent imaging obtains with the light field.

Our goal is to provide a model of coherent image formation that combines the utility of the light field with the comprehensive predictive power of the scalar field. The similarities between coherent and incoherent imaging motivate exploring how the scalar field and light field are related, which we address by synthesizing research across three different communities. Each community is concerned with a particular Fourier transform pair and has its own name for the light field. In optics, the pair is position and direction, and Walther discovered the first generalized radiance function by matching power predictions made with radiometry and scalar field theory [7]. In quantum physics, the pair is position and momentum, and Wigner discovered the first quasi-probability distribution, or phase-space distribution, as an aid to computing the expectation value of a quantum operator [8]. In signal processing, the pair is time and frequency, and while instantaneous spectra were used as early as 1890 by Sommerfeld, Ville is generally credited with discovering the first nontrivial quadratic time-frequency distribution by considering how to distribute the energy of a signal over time and frequency [9]. Walther, Wigner, and Ville independently arrived at essentially the same function, which is one of the ways to express a light field for coherent radiation in terms of the scalar field.

The light field has its roots in radiometry, a phenomenological theory of radiative power transport that began with Herschel’s observations of the sun [10], developed through the work of astrophysicists such as Chandrasekhar [11], and culminated with its grounding in electromagnetic field theory by Friberg et al. [12]. The light field represents radiance, which is the fundamental quantity in radiometry, defined as power per unit projected area per unit solid angle. Illuminating engineers would integrate radiance to compute power quantities, although no one could validate these calculations with the electromagnetic field theory formulated by Maxwell. Gershun was one of many physicists who attempted to physically justify radiometry, and who introduced the phrase *light field* to represent a three-dimensional vector field analogous to the electric and magnetic fields [13]. Gershun’s light field is a degenerate version of the one we discuss, and more closely resembles the time-averaged Poynting vector that appears in a rigorous derivation of geometric optics [14]. Subsequently, Walther generalized radiometry to coherent radiation in two different ways [7, 15], and Wolf connected Walther’s work to quantum physics [16], ultimately leading to the discovery of many more generalized radiance functions [17] and a firm foundation for radiometry [12].

Meanwhile, machine vision researchers desired a representation for all the possible pictures a pinhole camera might take in space-time, which led to the current formulation of the light field. Inspired by Leonardo da Vinci, Adelson and Bergen defined a plenoptic function to describe “everything that can be seen” as the intensity recorded by a pinhole camera, parametrized by position, direction, time, and wavelength [18]. Levoy and Hanrahan tied the plenoptic function more firmly to radiometry, by redefining Gershun’s phrase *light field* to mean radiance parametrized by position and direction [1]. Gortler et al. introduced the

same construct, but instead called it the lumigraph [19]. The *light field* is now the dominant terminology used in incoherent imaging contexts.

Our contribution is to describe and characterize all the ways to extend the light field to coherent radiation, and to interpret coherent image formation using the resulting extended light fields. We call our extended light fields *quasi light fields*, which are analogous to the generalized radiance functions of optics, the quasi-probability and phase-space distributions of quantum physics, and the quadratic class of time-frequency distributions of signal processing. Agarwal et al. have already extended the light field to coherent radiation [17], and the signal processing community has already classified all of the ways to distribute power over time and frequency [20]. Both have traced their roots to quantum physics. But to our knowledge, no one has connected the research to show 1) that the quasi light fields represent *all* the ways to extend the light field to coherent radiation, and 2) that the signal processing classification informs which quasi light field to use for a specific application. We further contextualize the references, making any unfamiliar literature more accessible to specialists in other areas.

Our paper is organized as follows. We describe the traditional light field in Section 2. We formulate the quasi light fields in Section 3, by reviewing and relating the relevant research in optics, quantum physics, and signal processing. In Section 4, we describe how to capture quasi light fields, discuss practical sampling issues, and illustrate the impact of light field choice on energy localization. In Section 5, we describe how to form images with quasi light fields. We derive a light field camera, demonstrate and compensate for diffraction limitations in the near zone, and generalize the classic beamforming algorithm in sensor array processing. We conclude the paper in Section 6, where we remark on the utility of quasi light fields and future perspectives on image formation.

2 The Traditional Light Field

The light field is a useful tool for incoherent imaging because it acts as an intermediary between the camera and the picture, decoupling information capture and image production: the camera measures the light field, from which many different traditional pictures can be computed. We define a pixel in the image of a scene by a surface patch σ and a virtual aperture (Figure 1). Specifically, we define the pixel value as the power P radiated by σ towards the aperture, just as an ideal single-lens camera would measure. According to radiometry, P is an integral over a bundle of light field rays [21]:

$$P = \int_{\sigma} \int_{\Omega_r} L(\mathbf{r}, \mathbf{s}) \cos \psi \, d^2s \, d^2r, \quad (1)$$

where $L(\mathbf{r}, \mathbf{s})$ is the radiance at position \mathbf{r} and in unit direction \mathbf{s} , ψ is the angle that \mathbf{s} makes with the surface normal at \mathbf{r} , and Ω_r is the solid angle subtended by the virtual aperture at \mathbf{r} . The images produced by many different conventional cameras can be computed from the light field using (1) [22].




Figure 1: We can compute the value of each pixel in an image produced by an arbitrary virtual camera, defined as the power emitted from a scene surface patch towards a virtual aperture, by integrating an appropriate bundle of light field rays that have been previously captured with remote hardware.

The light field has an important property that allows us to measure it remotely: the light field is constant along rays in a lossless medium [21]. To measure the light field on the surface of a scene, we follow the rays for the images we are interested in, and intercept those rays with our camera hardware (Figure 1). However, our hardware must be capable of measuring the radiance at a point and in a specific direction; a conventional camera that simply measures the irradiance at a point is insufficient. We can discern directional power flow using a lens array, as is done in a plenoptic camera [2].

In order to generate coherent images using the same framework described above, we must overcome three challenges. First, we must determine how to measure power flow by position and direction to formulate a coherent light field. Second, we must capture the coherent light field remotely and be able to infer behavior at the scene surface. Third, we must be able to use (1) to produce correct power values, so that we can form images by integrating over the coherent light field. We address each challenge in a subsequent section.

3 Formulating Quasi Light Fields

We motivate, systematically generate, and characterize the quasi light fields by relating existing research. We begin in Section 3.1 with research in optics that frames the challenge of extending the light field to coherent radiation in terms of satisfying a power constraint required for radiometry to make power predictions consistent with scalar field theory. While useful in developing an intuition for quasi light fields, the power constraint does not allow us to easily determine the quasi light fields. We therefore proceed in Section 3.2 to describe research in quantum physics that systematically generates quasi light fields satisfying the power constraint, and that shows how the quasi light fields are true extensions that reduce to the traditional light field under certain conditions. While useful for generating the quasi light fields, the quantum physics approach does not allow us to easily characterize them. Therefore, in Section 3.3 we map the generated quasi light fields to the quadratic class of

time-frequency distributions, which has been extensively characterized and classified by the signal processing community. By relating research in optics, quantum physics, and signal processing, we express all the ways to extend the light field to coherent radiation, and provide insight on how to select an appropriate quasi light field for a particular application.

We assume a perfectly coherent complex scalar field $U(\mathbf{r})$ at a fixed frequency ν for simplicity, although we comment in Section 6 on how to extend the results to broadband, partially coherent radiation. The radiometric theory we discuss assumes a planar source at $z = 0$. Consequently, although the light field is defined in three-dimensional space, much of our analysis is confined to planes $z = z_0$ parallel to the source. Therefore, for convenience, we use $\mathbf{r} = (x, y, z)$ and $\mathbf{s} = (s_x, s_y, s_z)$ to indicate three-dimensional vectors and $\mathbf{r}_\perp = (x, y)$ and $\mathbf{s}_\perp = (s_x, s_y)$ to indicate two-dimensional, projected versions.

3.1 Intuition from Optics

An extended light field must produce accurate power transport predictions consistent with rigorous theory; thus the power computed from the scalar field using wave optics determines the allowable light fields via the laws of radiometry. One way to find extended light fields is to guess a light field equation that satisfies this power constraint, which is how Walther identified the first extended light field [7]. The scenario involves a planar source at $z = 0$ described by $U(\mathbf{r})$, and a sphere of large radius ρ centered at the origin. We use scalar field theory to compute the flux through part of the sphere, and then use the definition of radiance to determine the light field from the flux.

According to scalar field theory, the differential flux $d\Phi$ through a portion of the sphere subtending differential solid angle $d\Omega$ is given by integrating the radial component of the energy flux density vector \mathbf{F} . From diffraction theory, the scalar field in the far zone is

$$U^\infty(\rho\mathbf{s}) = -\frac{2\pi i}{k} s_z \frac{\exp(ik\rho)}{\rho} a(\mathbf{s}) \quad (2)$$

where $k = 2\pi/\lambda$ is the wave number, λ is the wavelength, and

$$a(\mathbf{s}) = \left(\frac{k}{2\pi}\right)^2 \int U(\mathbf{r}) \exp(-ik\mathbf{s} \cdot \mathbf{r}) d^2r \quad (3)$$

is the plane wave component in direction \mathbf{s} [23]. Now

$$\mathbf{F}^\infty(\rho\mathbf{s}) = \left(\frac{2\pi}{k}\right)^2 a(\mathbf{s}) a^*(\mathbf{s}) \frac{s_z^2}{\rho^2} \mathbf{s}, \quad (4)$$

so that

$$d\Phi = \left(\frac{2\pi}{k}\right)^2 s_z^2 a(\mathbf{s}) a^*(\mathbf{s}) d\Omega. \quad (5)$$

According to radiometry, radiant intensity is flux per unit solid angle

$$I(\mathbf{s}) = \frac{d\Phi}{d\Omega} = \left(\frac{2\pi}{k}\right)^2 s_z^2 a(\mathbf{s}) a^*(\mathbf{s}). \quad (6)$$

Radiance is $I(\mathbf{s})$ per unit projected area [21], and this is where the guessing happens: there are many ways to distribute (6) over projected area by factoring out s_z and an outer integral over the source plane, but none yield light fields that satisfy all the traditional properties of radiance [24]. One way to factor (6) is to substitute the expression for $a(\mathbf{s})$ from (3) into (6) and change variables:

$$I(\mathbf{s}) = s_z \int \left[\left(\frac{k}{2\pi} \right)^2 s_z \int U \left(\mathbf{r} + \frac{1}{2} \mathbf{r}' \right) U^* \left(\mathbf{r} - \frac{1}{2} \mathbf{r}' \right) \exp(-ik\mathbf{s}_\perp \cdot \mathbf{r}'_\perp) d^2r' \right] d^2r. \quad (7)$$

The bracketed expression is Walther's first extended light field

$$L^W(\mathbf{r}, \mathbf{s}) = \left(\frac{k}{2\pi} \right)^2 s_z \mathfrak{W}(\mathbf{r}, \mathbf{s}_\perp/\lambda), \quad (8)$$

where

$$\mathfrak{W}(\mathbf{r}, \mathbf{s}_\perp) = \int U \left(\mathbf{r} + \frac{1}{2} \mathbf{r}' \right) U^* \left(\mathbf{r} - \frac{1}{2} \mathbf{r}' \right) \exp(-i2\pi\mathbf{s}_\perp \cdot \mathbf{r}'_\perp) d^2r' \quad (9)$$

is the Wigner distribution [25]. We may manually factor (6) differently to obtain other extended light fields in an *ad hoc* manner, but it is hard to find and verify the properties of all extended light fields this way, and we would have to individually analyze each light field that we do manage to find. So instead, we pursue a systematic approach to exhaustively identify and characterize the extended light fields that guarantee the correct radiant intensity in (6).

3.2 Generating Explicit Extensions from Quantum Physics

The mathematics of quantum physics provides us with a systematic extended light field generator that factors the radiant intensity in (6) in a structured way. Walther's extended light field in (8) provides the hint for this connection between radiometry and quantum physics. Specifically, Wolf recognized the similarity between Walther's light field and the Wigner phase-space distribution [8] from quantum physics [16]. Subsequently, Agarwal, Foley, and Wolf repurposed the mathematics behind phase-space representation theory to generate new light fields instead of distributions [17]. We summarize their approach, define the class of quasi light fields, describe how quasi light fields extend traditional radiometry, and show how quasi light fields can be conveniently expressed as filtered Wigner distributions.

Agarwal et al.'s key insight was to introduce a position operator $\hat{\mathbf{r}}_\perp$ and a direction operator $\hat{\mathbf{s}}_\perp$ that obey the commutation relations [26]

$$[\hat{x}, \hat{s}_x] = i\lambda/2\pi, \quad [\hat{y}, \hat{s}_y] = i\lambda/2\pi, \quad (10)$$

and to map the different ways of ordering the operators to different extended light fields. This formulation is valuable for two reasons. First, (10) is analogous to the quantum-mechanical relations for position and momentum, allowing us to exploit the phase-space distribution generator from quantum physics for our own purposes, thereby providing an explicit formula

for extended light fields. Second, in the geometric optics limit as $\lambda \rightarrow 0$, the operators commute per (10), so that all of the extended light fields collapse to the same function that can be related to the traditional light field. Therefore, Agarwal et al.'s formulation not only provides us with different ways of expressing the light field for coherent radiation, but also explains how these differences arise as the wavelength becomes non-negligible.

We now summarize the phase-space representation calculus that Agarwal and Wolf invented [27] to map operator orderings to functions, which Agarwal et al. later applied to radiometry [17], culminating in a formula for extended light fields. The phase-space representation theory generates a function \tilde{L}^Ω from any operator \hat{L} for each distinct way Ω of ordering collections of $\hat{\mathbf{r}}_\perp$ and $\hat{\mathbf{s}}_\perp$. So by choosing a specific \hat{L} defined by its matrix elements using the Dirac notation [26]

$$\langle \mathbf{r}_\perp^R | \hat{L} | \mathbf{r}_\perp^C \rangle = U(\mathbf{r}^R) U^*(\mathbf{r}^C), \quad (11)$$

and supplying \hat{L} as input, we obtain the extended light fields

$$L^\Omega(\mathbf{r}, \mathbf{s}) = \left(\frac{k}{2\pi} \right)^2 s_z \tilde{L}^\Omega(\mathbf{r}_\perp, \mathbf{s}_\perp) \quad (12)$$

as outputs. The power constraint from Section 3.1 translates to a minor constraint on the allowed orderings Ω , so that L^Ω can be factored from (6). Finally, there is an explicit formula for L^Ω [27], which in Friberg et al.'s form [12] reads

$$\begin{aligned} L^\Omega(\mathbf{r}, \mathbf{s}) &= \frac{k^2}{(2\pi)^4} s_z \iiint \tilde{\Omega}(\mathbf{u}, k\mathbf{r}_\perp'') \exp[-i\mathbf{u} \cdot (\mathbf{r}_\perp - \mathbf{r}_\perp')] \exp(-iks_\perp \cdot \mathbf{r}_\perp'') \\ &\quad \times U\left(\mathbf{r}' + \frac{1}{2}\mathbf{r}''\right) U^*\left(\mathbf{r}' - \frac{1}{2}\mathbf{r}''\right) d^2u d^2r' d^2r'', \end{aligned} \quad (13)$$

where $\tilde{\Omega}$ is a functional representation of the ordering Ω .

Previous research has related the extended light fields L^Ω to the traditional light field, by examining how the L^Ω behave for globally incoherent light of a small wavelength, an environment technically modeled by a quasi-homogeneous source in the geometric optics limit where $\lambda \rightarrow 0$. As $\lambda \rightarrow 0$, $\hat{\mathbf{r}}_\perp$ and $\hat{\mathbf{s}}_\perp$ commute per (10), so that all orderings Ω are equivalent and all of the extended light fields L^Ω collapse to the same function. Since, in the source plane, Foley and Wolf showed that one of those light fields behaves like traditional radiance [28] for globally incoherent light of a small wavelength, all of the L^Ω behave like traditional radiance for globally incoherent light of a small wavelength. Furthermore, Friberg et al. showed that many of the L^Ω are constant along rays, for globally incoherent light of a small wavelength [12]. The L^Ω thereby subsume the traditional light field, and *globally incoherent light of a small wavelength* is the environment in which traditional radiometry holds.

To more easily relate L^Ω to the signal processing literature, we conveniently express L^Ω as a filtered Wigner distribution. We introduce a function Π and substitute

$$\tilde{\Omega}(\mathbf{u}, \mathbf{v}) = \iint \Pi(-\mathbf{a}, -\mathbf{b}) \exp[-i(\mathbf{a} \cdot \mathbf{u} + \mathbf{b} \cdot \mathbf{v})] d^2a d^2b \quad (14)$$

into (13), integrate first over \mathbf{u} , then over \mathbf{a} , and finally substitute $\mathbf{b} = \mathbf{s}'_{\perp} - \mathbf{s}_{\perp}$:

$$\begin{aligned} L^{\Omega}(\mathbf{r}, \mathbf{s}) &= \left(\frac{k}{2\pi}\right)^2 s_z \iint \Pi(\mathbf{r}_{\perp} - \mathbf{r}'_{\perp}, \mathbf{s}_{\perp} - \mathbf{s}'_{\perp}) \mathfrak{W}(\mathbf{r}', \mathbf{s}'_{\perp}/\lambda) d^2 r' d^2 s' \\ &= \left(\frac{k}{2\pi}\right)^2 s_z \Pi(\mathbf{r}_{\perp}, \mathbf{s}_{\perp}) \otimes \mathfrak{W}(\mathbf{r}, \mathbf{s}_{\perp}/\lambda). \end{aligned} \quad (15)$$

The symbol \otimes in (15) denotes convolution in both \mathbf{r}_{\perp} and \mathbf{s}_{\perp} . Each filter kernel Π yields a different light field. There are only minor restrictions on Π , or equivalently on $\tilde{\Omega}$. Specifically, Agarwal and Wolf's calculus requires that [27]

$$1/\tilde{\Omega} \text{ is an entire analytic function with no zeros on the real component axes.} \quad (16)$$

Agarwal et al.'s derivation additionally requires that

$$\tilde{\Omega}(\mathbf{0}, \mathbf{v}) = 1 \text{ for all } \mathbf{v}, \quad (17)$$

so that L^{Ω} satisfies the laws of radiometry and is consistent with (6) [17].

We call the functions L^{Ω} , the restricted class of extended light fields that we have systematically generated, *quasi light fields*, in recognition of their connection with quasi-probability distributions in quantum physics.

3.3 Characterization from Signal Processing

Although we have identified the quasi light fields and justified how they extend the traditional light field, we must still show that we have found all possible ways to extend the light field to coherent radiation, and we must indicate how to select a quasi light field for a specific application. We address both concerns by relating quasi light fields to bilinear forms of U and U^* that are parameterized by position and direction. First, such bilinear forms reflect all the different ways to represent the energy distribution of a complex signal in signal processing, and therefore contain all possible extended light fields, allowing us to identify any unaccounted for by quasi light fields. Second, we may use the signal processing classification of bilinear forms to characterize quasi light fields and guide the selection of one for an application.

To relate quasi light fields to bilinear forms, we must express the filtered Wigner distribution in (15) as a bilinear form. Towards this end, we first express the filter kernel Π in terms of another function K :

$$\Pi(\mathbf{a}, \mathbf{b}) = \int K \left(-\mathbf{a} + \frac{\lambda}{2}\mathbf{v}, -\mathbf{a} - \frac{\lambda}{2}\mathbf{v} \right) \exp(-i2\pi\mathbf{b} \cdot \mathbf{v}) d^2 v. \quad (18)$$

We substitute (18) into (15), integrate first over \mathbf{s}'_{\perp} , then over \mathbf{v} , and finally substitute

$$\mathbf{r}^R = \mathbf{r}' + \frac{1}{2}\mathbf{r}'', \quad \mathbf{r}^C = \mathbf{r}' - \frac{1}{2}\mathbf{r}'' \quad (19)$$

to express the quasi light field as

$$L(\mathbf{r}, \mathbf{s}) = \left(\frac{k}{2\pi}\right)^2 s_z \iint U(\mathbf{r}^R) \left\{ K(\mathbf{r}_\perp^R - \mathbf{r}_\perp, \mathbf{r}_\perp^C - \mathbf{r}_\perp) \right. \\ \left. \times \exp[-ik\mathbf{s}_\perp \cdot (\mathbf{r}_\perp^R - \mathbf{r}_\perp^C)] \right\} U^*(\mathbf{r}^C) d^2r^R d^2r^C. \quad (20)$$

We recognize that (20) is a bilinear form of U and U^* , with kernel indicated by the braces.

The structure of the kernel of the bilinear form in (20) limits L to a shift-invariant energy distribution. Specifically, translating the scalar field in (20) in position and direction orthogonal to the z -axis according to

$$U(\mathbf{r}) \rightarrow U(\mathbf{r} - \mathbf{r}^0) \exp(ik\mathbf{s}_\perp^0 \cdot \mathbf{r}_\perp) \quad (21)$$

results in a corresponding translation in position and direction in the light field, after rearranging terms:

$$L(\mathbf{r}, \mathbf{s}) \rightarrow L(\mathbf{r} - \mathbf{r}^0, \mathbf{s} - \mathbf{s}^0). \quad (22)$$

Such shift-invariant bilinear forms comprise the *quadratic class* of time-frequency distributions, which is sometimes misleadingly referred to as Cohen's class [20].

The quasi light fields represent *all* possible ways of extending the light field to coherent radiation. This is because any reasonably defined extended light field must be shift-invariant in position and direction, as translating and rotating coordinates should modify the scalar field and light field representations in corresponding ways. Thus, on the one hand, an extended light field must be a quadratic time-frequency distribution. On the other hand, (20) implies that quasi light fields span the entire class of quadratic time-frequency distributions, apart from the constraints on Π described at the end of Section 3.2. The constraint in (17) is necessary to satisfy the power constraint in (6), which any extended light field must satisfy. The remaining constraints in (16) are technical details concerning analyticity and the location of zeros; extended light fields strictly need not satisfy these mild constraints, but the light fields that are ruled out are well-approximated by light fields that satisfy them.

We obtain a concrete sensor array processing interpretation of quasi light fields by grouping the exponentials in (20) with U instead of K :

$$L(\mathbf{r}, \mathbf{s}) = \left(\frac{k}{2\pi}\right)^2 s_z \iint \left\{ U(\mathbf{r}^R) \exp[ik\mathbf{s} \cdot (\mathbf{r} - \mathbf{r}^R)] \right\} K(\mathbf{r}_\perp^R - \mathbf{r}_\perp, \mathbf{r}_\perp^C - \mathbf{r}_\perp) \\ \times \left\{ U(\mathbf{r}^C) \exp[ik\mathbf{s} \cdot (\mathbf{r} - \mathbf{r}^C)] \right\}^* d^2r^R d^2r^C. \quad (23)$$

The integral in (23) is the expected value of the energy of the output of a spatial filter with impulse response $\exp(ik\mathbf{s} \cdot \mathbf{r})$ applied to the scalar field, when using K to estimate the correlation $E[U(\mathbf{r}^R)U^*(\mathbf{r}^C)]$ by

$$U(\mathbf{r}^R) K(\mathbf{r}_\perp^R - \mathbf{r}_\perp, \mathbf{r}_\perp^C - \mathbf{r}_\perp) U^*(\mathbf{r}^C). \quad (24)$$

That is, the choice of quasi light field corresponds to a choice of how to infer coherence structure from scalar field measurements. In adaptive beamforming, the spatial filter $\exp(ik\mathbf{s} \cdot$

\mathbf{r}) focuses a sensor array on a particular plane wave component, and K serves a similar role as the covariance matrix taper that gives rise to design features such as diagonal loading [29]. But for our purposes, the sensor array processing interpretation in (23) allows us to cleanly separate the choice of quasi light field in K from the plane wave focusing in the exponentials.

Several signal processing books meticulously classify the quadratic class of time-frequency distributions by their properties, and discuss distribution design and use for various applications [25, 20]. We can use these resources to design quasi light fields for specific applications. For example, if we desire a light field with fine directional localization, we may first try the Wigner quasi light field in (8), which is a popular starting choice. We may then discover that we have too many artifacts from interfering spatial frequencies, called *cross terms*, and therefore wish to consider a reduced interference quasi light field. We might try the modified B-distribution, which is a particular reduced interference quasi light field that has a tunable parameter to suppress interference. Or, we may decide to design our own quasi light field in a transformed domain using ambiguity functions. The resulting tradeoffs can be tailoring to specific application requirements.

4 Capturing Quasi Light Fields

To capture an arbitrary quasi light field, we sample and process the scalar field. In incoherent imaging, traditional light fields are typically captured by instead making intensity measurements at a discrete set of positions and directions, as in done in the plenoptic camera [2]. While it is possible to apply the same technique to coherent imaging, only a small subset of quasi light fields that exhibit poor localization properties can be captured this way. In comparison, all quasi light fields can be computed from the scalar field, as in (15). We therefore sample the scalar field with a discrete set of sensors placed at different positions in space, and subsequently process the scalar field measurements to compute the desired quasi light field. We describe the capture process for three specific quasi light fields in Section 4.1, and demonstrate the different localization properties of these quasi light fields via simulation in Section 4.2.

4.1 Sampling the Scalar Field

To make the capture process concrete, we capture three different quasi light fields. For simplicity, we consider a two-dimensional scene and sample the scalar field with a linear array of sensors regularly spaced along the y -axis (Figure 2). With this geometry, the scalar field U is parameterized by a single position variable y , and the discrete light field ℓ is parameterized by y and the direction component s_y . The sensor spacing is $d/2$, which we assume is fine enough to ignore aliasing effects. This assumption is practical for long-wavelength applications such as millimeter-wave radar. For other applications, aliasing can be avoided by applying an appropriate pre-filter. From the sensor measurements, we compute three different quasi light fields, including the spectrogram and the Wigner.

Although the spectrogram quasi light field is attractive because it can be captured like a




Figure 2: We capture a discrete quasi light field ℓ by sampling the scalar field at regularly-spaced sensors and processing the resulting measurements. We may optionally apply an aperture stop T to mimic traditional light field capture, but this restricts us to capturing quasi light fields with poor localization properties.

traditional light field by making intensity measurements, it exhibits poor localization properties. Zhang and Levoy explain [30] how to capture the spectrogram by placing an aperture stop specified by a transmission function T over the desired position y before computing a Fourier transform to extract the plane wave component in the desired direction s_y , and previously Ziegler et al. used the spectrogram as a coherent light field to represent a hologram [4]. The spectrogram is an important quasi light field because it is the building block for the quasi light fields that can be directly captured by making intensity measurements, since all non-negative quadratic time-frequency distributions, and therefore all non-negative quasi light fields, are sums of spectrograms [20]. Ignoring constants and s_z , we compute the discrete spectrogram from the scalar field samples by

$$\ell^S(y, s_y) = \left| \sum_n T(nd)U(y + nd) \exp(-iknds_y) \right|^2. \quad (25)$$

The Wigner quasi light field is a popular choice that exhibits good energy localization in position and direction [20]. We already identified the Wigner quasi light field in (8); the discrete version is

$$\ell^W(y, s_y) = \sum_n U(y + nd/2)U^*(y - nd/2) \exp(-iknds_y). \quad (26)$$

Evidently, the spectrogram and Wigner distribute energy over position and direction in very different ways. Per (25), the spectrogram first uses a Fourier transform to extract directional information and then computes a quadratic energy quantity, while the Wigner does the reverse, per (26). On the one hand, this reversal allows the Wigner to better localize energy in position and direction, since the Wigner is not bound by the classical Fourier uncertainty principle as the spectrogram is. On the other hand, the Wigner’s nonlinearities introduce cross-term artifacts by coupling energy in different directions, thereby replacing the simple uncertainty principle with a more complicated set of tradeoffs [20].

We now introduce a third quasi light field to capture, in order to help us understand the implications of requiring quasi light fields to exhibit traditional light field properties. Specifically, traditional light fields have real non-negative values that are zero where the scalar field is zero, whereas no quasi light field behaves this way [24]. Although the spectrogram has non-negative values, the support of both the spectrogram and Wigner spills over into regions where the scalar field is zero. In contrast, the conjugate Rihaczek quasi light field, which can be obtained by substituting (3) for $a^*(\mathbf{s})$ in (6) and factoring, is identically zero at all positions where the scalar field is zero and for all directions in which the plane wave component is zero:

$$L^R(\mathbf{r}, \mathbf{s}) = s_z U^*(\mathbf{r}) \exp(ik\mathbf{s} \cdot \mathbf{r}) a(\mathbf{s}). \quad (27)$$

However, unlike the non-negative spectrogram and the real Wigner, the Rihaczek is complex-valued, as each of its discoverers independently observed: Walther in optics [15], Kirkwood in quantum physics [31], and Rihaczek in signal processing [32]. The discrete conjugate Rihaczek quasi light field is

$$\ell^R(y, s_y) = U^*(y) \exp(ikys_y) \sum_n U(nd) \exp(-iknds_y). \quad (28)$$

4.2 Localization Tradeoffs

Different quasi light fields localize energy in position and direction in different ways, so that the choice of quasi light field impacts the potential resolution achieved in an imaging application. We illustrate the diversity of behavior by simulating a plane wave propagating past a screen edge and computing the spectrogram, Wigner, and Rihaczek quasi light fields from scalar field samples (Figure 3). This simple scenario stresses the main tension between localization in position and direction: each quasi light field must encode the position of the screen edge as well as the downward direction of the plane wave. The quasi light fields serve as intermediate representations used to jointly estimate the position of the screen edge and the orientation of the plane wave.

Our simulation accurately models diffraction using our implementation of the angular spectrum propagation method, which is the same technique used in commercial optics software to accurately simulate wave propagation [33]. We propagate a plane wave with wavelength $\lambda = 3$ mm a distance $R = 50$ m past the screen edge, where we measure the scalar field and compute the three discrete light fields using (25), (26), and (28). To compute the light fields, we set $d = \lambda/10$, run the summations over $|n| \leq 10/\lambda$, and use a rectangular window function of width 10 cm for T . We plot ℓ^S , $|\ell^W|$, and $|\ell^R|$ in terms of the two-plane parameterization of the light field [1], so that each ray is directed from a point u in the plane of the screen towards a point y in the measurement plane, and so that $s_y = (y - u)/[R^2 + (y - u)^2]^{1/2}$.

We compare each light field's ability to estimate the position of the screen edge and the orientation of the plane wave (Figure 3). Geometric optics provides an ideal estimate: we should ideally only see rays pointing straight down ($u = y$) past the screen edge, corresponding to a diagonal line in the upper-right quadrant of the light field plots. Instead, we see

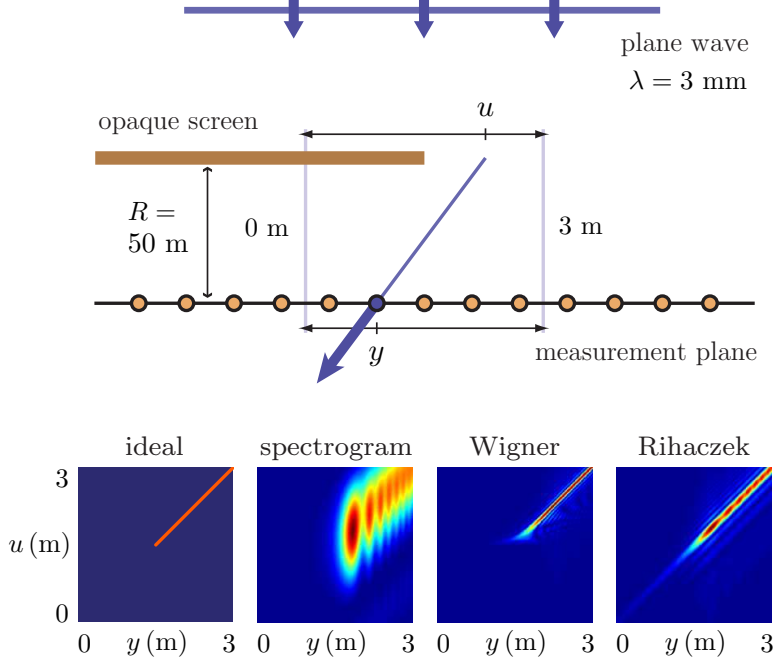


Figure 3: The spectrogram does not resolve a plane wave propagating past the edge of an opaque screen as well as other quasi light fields, such as the Wigner and Rihaczek. We capture all three quasi light fields by sampling the scalar field with sensors and processing the measurements according to (25), (26), and (28). The ringing and blurring in the light field plots indicate the diffraction fringes and energy localization limitations.

blurred lines with ringing. The ringing is physically accurate and indicates the diffraction fringes formed on the measurement plane. The blurring indicates localization limitations. While the spectrogram's window T can be chosen to narrowly localize energy in either position or direction, the Wigner narrowly localizes energy in both, depicting instantaneous frequency without being limited by the classical Fourier uncertainty principle [20].

It may seem that the Wigner light field is preferable to the others and the clear choice for all applications. While the Wigner light field possesses excellent localization properties, it exhibits cross-term artifacts due to interference from different plane wave components. An alternative quasi light field such as the Rihaczek can strike a balance between localization and cross-term artifacts, and therefore may be a more appropriate choice, as discussed at the end of Section 3.3. If our goal were to only estimate the position of the screen edge, we might prefer the spectrogram; to jointly estimate both position and plane wave orientation, we prefer the Wigner; and if there were two plane waves instead of one, we might prefer the Rihaczek. One thing is certain, however: we must abandon non-negative quasi light fields to achieve better localization tradeoffs, as all non-negative quadratic time-frequency distributions are sums of spectrograms and hence exhibit poor localization tradeoffs [20].

5 Image Formation

We wish to form images from quasi light fields for coherent applications similarly to how we form images from traditional light fields for incoherent applications, by using (1) to integrate bundles of light field rays to compute pixel values (Figure 1). However, simply selecting a particular captured quasi light field L and evaluating (1) raises three questions about the validity of the resulting image. First, is it meaningful to distribute coherent energy over surface area by factoring radiant intensity in (6)? Second, does the far-zone assumption implicit in radiometry and formalized in (2) limit the applicability of quasi field fields? And third, how do we capture quasi light field rays remotely if, unlike the traditional light field, quasi light fields need not be constant along rays?

The first question is a semantic one. For incoherent light of a small wavelength, we *define* an image in terms of the power radiating from a scene surface towards an aperture, and physics tells us that this uniquely specifies the image (Section 3), which may be expressed in terms of the traditional light field. If we attempt to generalize the same definition of an image to partially coherent, broadband light, and specifically to coherent light at a non-zero wavelength, we must ask how to isolate the power from a surface patch towards the aperture, according to classical wave optics. But there is no unique answer; different isolation techniques correspond to different quasi light fields. Therefore, to be well-defined, we must extend the definition of an image for coherent light to include a particular choice of quasi light field, which corresponds to a particular factorization of radiant intensity.

The second and third questions speak of assumptions in the formulation of quasi light fields and in the image formation from quasi light fields, which can lead to coherent imaging inaccuracies when these assumptions are not valid. Specifically, unless the scene surface and aperture are far apart, the far-zone assumption in (2) does not hold, so that quasi light fields are incapable of modeling near-zone behavior. Also, unless we choose a quasi light field that is constant along rays, such as an angle-impact Wigner function [34], remote measurements might not accurately reflect the light field at the scene surface [35], resulting in imaging inaccuracies. Therefore, in general, integrating bundles of remotely captured quasi light field rays produces an approximation of the image we have defined. We assess this approximation by building an accurate near-zone model in Section 5.1, simulating imaging performance of several coherent cameras in Section 5.2, and showing how our image formation procedure generalizes the classic beamforming algorithm in Section 5.3.

5.1 Near-Zone Radiometry

We take a new approach to formulating light fields for coherent radiation that avoids making the assumptions that 1) the measurement plane is far from the scene surface and 2) light fields are constant along rays. The resulting light fields are accurate in the near zone, and may be compared with quasi light fields to understand quasi light field limitations. The key idea is to express a near-zone light field $L(\mathbf{r}, \mathbf{s})$ on the measurement plane in terms of the infinitesimal flux at the point where the line containing the ray (\mathbf{r}, \mathbf{s}) intersects the scene surface (Figure 4). First we compute the scalar field at the scene surface, next we compute

the infinitesimal flux, and then we identify a light field that predicts the same flux using the laws of radiometry. In contrast with Walther’s approach (Section 3.1), 1) we do not make the far-zone approximation as in (2), and 2) we formulate the light field in the measurement plane instead of in the source plane at the scene surface. Therefore, in forming an image from a near-zone light field, we are not limited to the far zone and we need not relate the light field at the measurement plane to the light field at the scene surface.




Figure 4: To ensure that integrating bundles of remote light field rays in the near zone results in an accurate image, we derive a light field $L_\rho^R(\mathbf{r}, \mathbf{s})$ in the measurement plane from the infinitesimal flux $d\Phi$ at the point \mathbf{r}^P where the ray originates from the scene surface patch. We thereby avoid making the assumptions that the measurement plane is far from the scene and that the light field is constant along rays.

The first step in deriving a near-zone light field L for the ray (\mathbf{r}, \mathbf{s}) is to use the scalar field on the measurement plane to compute the scalar field at the point \mathbf{r}^P where the line containing the ray intersects the scene surface. We choose coordinates so that the measurement plane is the xy -plane, the scene lies many wavelengths away in the negative $z < 0$ half-space, and \mathbf{r} is at the origin. We denote the distance between the source \mathbf{r}^P on the scene surface and the point of observation \mathbf{r} by ρ . Under a reasonable bandwidth assumption, the inverse diffraction formula expresses the scalar field at \mathbf{r}^P in terms of the scalar field on the measurement plane [36]:

$$U(\mathbf{r}^P) = \frac{ik}{2\pi} \int U(\mathbf{r}^M) \frac{-z^P}{|\mathbf{r}^P - \mathbf{r}^M|} \frac{\exp(-ik|\mathbf{r}^P - \mathbf{r}^M|)}{|\mathbf{r}^P - \mathbf{r}^M|} d^2r^M. \quad (29)$$

Next, we compute the differential flux $d\Phi$ through a portion of a sphere at \mathbf{r}^P subtending differential solid angle $d\Omega$. We obtain $d\Phi$ by integrating the radial component of the energy flux density vector

$$\mathbf{F}(\mathbf{r}^P) = -\frac{1}{4\pi k\nu} \left[\frac{\partial U^*}{\partial t} \nabla U + \frac{\partial U}{\partial t} \nabla U^* \right]. \quad (30)$$

To keep the calculation simple, we ignore amplitude decay across the measurement plane, approximating

$$|\mathbf{r}^P - \mathbf{r}^M| \approx |\mathbf{r}^P| \quad (31)$$

outside the exponential in (29), and

$$\frac{\partial}{\partial |\mathbf{r}^P|} |\mathbf{r}^P - \mathbf{r}^M| \approx 1, \quad (32)$$

when evaluating (30), resulting in

$$\mathbf{F}(-\rho\mathbf{s}) = \left(\frac{2\pi}{k}\right)^2 \tilde{a}(-\rho\mathbf{s}) \tilde{a}^*(-\rho\mathbf{s}) \frac{s_z^2}{\rho^2} \mathbf{s}, \quad (33)$$

where

$$\tilde{a}(-\rho\mathbf{s}) = \left(\frac{k}{2\pi}\right)^2 \int U(\mathbf{r}^M) \exp(-ik|\mathbf{r} - \rho\mathbf{s} - \mathbf{r}^M|) d^2r^M. \quad (34)$$

Thus,

$$d\Phi = \left(\frac{2\pi}{k}\right)^2 s_z^2 \tilde{a}(-\rho\mathbf{s}) \tilde{a}^*(-\rho\mathbf{s}) d\Omega. \quad (35)$$

Finally, we factor out s_z and an outer integral over surface area from $d\Phi/d\Omega$ to determine a near-zone light field. Unlike in Section 3.1, the nonlinear exponential argument in \tilde{a} complicates the factoring. Nonetheless, we obtain a near-zone light field that generalizes the Rihaczek by substituting (34) for \tilde{a}^* in (35). After factoring and freeing \mathbf{r} from the origin by substituting $\mathbf{r} - \rho\mathbf{s}$ for $-\rho\mathbf{s}$, we obtain

$$\begin{aligned} L_\rho^R(\mathbf{r}, \mathbf{s}) &= s_z U^*(\mathbf{r}) \exp(ik\rho) \tilde{a}(\mathbf{r} - \rho\mathbf{s}) \\ &= \left(\frac{k}{2\pi}\right)^2 s_z U^*(\mathbf{r}) \exp(ik\rho) \int U(\mathbf{r}^M) \exp(-ik|\mathbf{r} - \rho\mathbf{s} - \mathbf{r}^M|) d^2r^M, \end{aligned} \quad (36)$$

where the subscript ρ reminds us of this near-zone light field's dependence on distance.

L_ρ^R is evidently neither the traditional light field nor a quasi light field, as it depends directly on the scene geometry through an additional distance parameter. This distance parameter ρ is a function of \mathbf{r} , \mathbf{s} , and the geometry of the scene; it is the distance along \mathbf{s} between the scene surface and \mathbf{r} . We may integrate L_ρ^R over a bundle of rays to compute the image pixel values just like any other light field, as long as we supply the right value of ρ for each ray. In contrast, quasi light fields are incapable of modeling optical propagation in the near zone, as it is insufficient to specify power flow along rays. We must also know the distance between the source and point of measurement along each ray.

We can obtain near-zone generalizations of all quasi light fields through the sensor array processing interpretation in Section 3.3. Recall that each quasi light field corresponds to a particular choice of the function K in (23). For example, setting $K(\mathbf{a}, \mathbf{b}) = \delta(\mathbf{b})$, where δ is the Dirac delta function, yields the Rihaczek quasi light field L^R in (27). To generalize quasi light fields to the near zone, we focus at a point instead of a plane wave component by using a spatial filter with impulse response $\exp(-ik|\mathbf{r} - \rho\mathbf{s}|)$ instead of $\exp(iks \cdot \mathbf{r})$ in (23). Then, choosing $K(\mathbf{a}, \mathbf{b}) = \delta(\mathbf{b})$ yields L_ρ^R , the near-zone generalization of the Rihaczek in (36), and choosing other functions K yield near-zone generalizations of the other quasi light fields.

5.2 Near-Zone Diffraction Limitations

We compute and compare image pixel values using the Rihaczek quasi light field L^R and its near-zone generalization L_ρ^R , demonstrating how all quasi light fields implicitly make the Fraunhofer diffraction approximation that limits accurate imaging to the far zone. First, we construct coherent cameras from L^R and L_ρ^R . For simplicity, we consider a two-dimensional scene and sample the light fields, approximating the integral over a bundle of rays (Figure 1) by the summation of discrete rays directed from the center \mathbf{r}^P of the scene surface patch to each sensor on a virtual aperture of diameter A , equally spaced every distance d in the measurement plane (Figure 5a). Ignoring constants and s_z , we compute the pixel values for a far-zone camera from the Rihaczek quasi light field in (27),

$$P^R = \sum_{|nd| < A/2} [U(nd) \exp(-iknds_y^n)]^* \left[\sum_m U(md) \exp(-ikmds_y^n) \right], \quad (37)$$

and for a near-zone camera from the near-zone generalization of the Rihaczek in (36),

$$P_\rho^R = \left[\sum_{|nd| < A/2} U(nd) \exp(-ik\Delta_n) \right]^* \left[\sum_m U(md) \exp(-ik\Delta_m) \right]. \quad (38)$$

In (37), \mathbf{s}^n denotes the unit direction from \mathbf{r}^P to the n^{th} sensor, and in (38), Δ_n denotes the distance between \mathbf{r}^P and the n^{th} sensor.

By comparing the exponentials in (37) with those in (38), we see that the near-zone camera aligns the sensor measurements along spherical wavefronts diverging from the point of focus \mathbf{r}^P , while the far-zone camera aligns measurements along plane wavefront approximations (Figure 5b). Spherical wavefront alignment makes physical sense in accordance with the Huygens-Fresnel principle of diffraction, while approximating spherical wavefronts with plane wavefronts is reminiscent of Fraunhofer diffraction. In fact, the far-zone approximation in (2) used to derive quasi light fields follows directly from the Rayleigh-Sommerfeld diffraction integral by linearizing the exponentials, which is precisely Fraunhofer diffraction. Therefore, all quasi light fields are only valid for small Fresnel numbers, when the source and point of measurement are sufficiently far away from each other.

We expect the near-zone camera to outperform the far-zone camera in near-zone imaging applications, which we demonstrate by comparing their ability to resolve small targets moving past their field of view. As a baseline, we introduce a third camera with non-negative pixel values P_ρ^B by restricting the summation over m in (38) to $|md| < A/2$, which results in the beamformer camera used in sensor array processing [5, 37]. Alternatively, we could extend the summation over n in (38) to the entire array, but this would average anisotropic responses over a wider aperture diameter, resulting in a different image. We simulate an opaque screen containing a pinhole that is backlit with a coherent plane wave (Figure 6). The sensor array is $D = 2$ m wide and just $R = 1$ m away from the screen. The virtual aperture is $A = 10$ cm wide and the camera is focused on a fixed 1 mm pixel straight ahead on the screen. The pinhole has width 1 mm, which is smaller than the wavelength $\lambda = 3$




Figure 5: The near-zone light field results in a camera that aligns spherical wavefronts diverging from the point of focus \mathbf{r}^P , in accordance with the Huygens-Fresnel principle of diffraction, while quasi light fields result in cameras that align plane wavefront approximations, in accordance with Fraunhofer diffraction. Quasi light fields are therefore only accurate in the far zone. We derive both cameras by approximating the integral over a bundle of rays by the summation of discrete light field rays (a), and we interpret the operation of each camera by how they align sensor measurements along wavefronts from \mathbf{r}^P (b).

mm, so the plane wavefronts bend into slightly spherical shapes via diffraction. We move the pinhole to the right, recording pixel values $|P^R|$, $|P_\rho^R|$, and P_ρ^B for each camera at each pinhole position. Due to the nature of the coherent combination of the sensor measurements that produces the pixel values, each camera records a multi-lobed response. The width of the main lobe indicates the near-zone resolution of the camera.

The near-zone camera is able to resolve the pinhole down to its actual size of 1 mm, greatly outperforming the far-zone camera which records a blur 66 cm wide, and even outperforming the beamformer camera. Neither comparison is surprising. First, with a Fresnel number of $D^2/R\lambda \approx 1333$, the Fraunhofer approximation implicitly made by quasi light fields does not hold for this scenario, so we expect the far-zone camera to exhibit poor resolution. Second, the near-zone camera uses the entire $D = 2$ m array instead of just the sensors on the virtual aperture that the beamformer camera is restricted to, and the extra sensors lead to improved resolution.




Figure 6: Images of nearby objects formed from pure quasi light fields are blurry. In the scene, a small backlit pinhole moves across the field of view of a sensor array that implements three cameras, each computing one pixel value for each pinhole position, corresponding to a fixed surface patch. As the pinhole crosses the fixed scene surface patch, the near-zone camera resolves the pinhole down to its actual size of 1 mm, while the far-zone camera records a blur 66 cm wide.

5.3 Generalized Beamforming

We compare image formation from light fields with traditional perspectives on coherent image formation, by relating quasi light fields and our coherent cameras with the classic beamforming algorithm used in many coherent imaging applications, including ultrasound [5] and radar [37]. The beamforming algorithm estimates a spherical wave diverging from a point of focus \mathbf{r}^P by delaying and averaging sensor measurements. When the radiation is narrowband, the delays are approximated by phase shifts. With the sensor array geometry from Section 5.2, the beamformer output is

$$g = \sum_m T(md)U(md) \exp(-ik\Delta_m), \quad (39)$$

where the $T(md)$ are amplitude weights used to adjust the beamformer's performance. As \mathbf{r}^P moves into the far zone,

$$\Delta_m - \Delta_0 \rightarrow mds_y^m \rightarrow mds_y^0, \quad (40)$$

so that apart from a constant phase offset, (39) becomes a short-time Fourier transform

$$g^\infty = \sum_m T(md)U(md) \exp(-ikmds_y^0). \quad (41)$$

Evidently, $|g^\infty|^2$ is a spectrogram quasi light field, and we may select T to be a narrow window about a point \mathbf{r} to capture $L^S(\mathbf{r}, \mathbf{s}^0)$. We have already seen how quasi light fields generalize the spectrogram.

Beamformer applications instead typically select T to be a wide window to match the desired virtual aperture, and assign the corresponding pixel value to the output power $|g|^2$. We can decompose the three cameras in Section 5.2 into such beamformers. First, we write P_ρ^R in (38) in terms of two different beamformers,

$$P_\rho^R = g_1^* g_2, \quad (42)$$

where

$$g_1 = \sum_{|nd| < A/2} U(nd) \exp(-ik\Delta_n) \quad (43)$$

and

$$g_2 = \sum_m U(md) \exp(-ik\Delta_m), \quad (44)$$

so that the windows for g_1 and g_2 are rectangular with widths matching the aperture A and sensor array D , respectively. Next, by construction

$$P_\rho^B = |g_1|^2. \quad (45)$$

Finally, in the far zone, $\mathbf{s}^n \rightarrow \mathbf{s}^0$ in (37) so that

$$P^R \rightarrow (g_1^\infty)^* g_2^\infty, \quad (46)$$

where g_1^∞ and g_2^∞ are given by (41) with the windows T used in (43) and (44). In other words, the near-zone camera is the Hermitian product of two different beamformers, and is equivalent to the far-zone camera in the far zone.

We interpret the role of each component beamformer from the derivation of (38). Beamformer g_1^* aggregates power contributions across the aperture using measurements of the conjugate field U^* on the aperture, while beamformer g_2 isolates power from the point of focus using all available measurements of the field U . In this manner, the tasks of aggregating and isolating power contributions are cleanly divided between the two beamformers, and each beamformer uses the measurements from those sensors appropriate to its task. In contrast, the beamformer camera uses the same set of sensors for both the power aggregation and isolation tasks, thereby limiting its ability to optimize over both tasks.

The near-zone camera achieves a new tradeoff between resolution and anisotropic sensitivity. We noted that the near-zone camera exhibits better resolution than the beamformer, for the same virtual aperture (Figure 6). This is not an entirely fair comparison because the

near-zone camera is using sensor measurements outside the aperture, and indeed, a beamformer using the entire array would achieve comparable resolution. However, extending the aperture to the entire array results in a different image, as anisotropic responses are averaged over a wider aperture diameter. We interpret the near-zone camera’s behavior by computing the magnitude

$$|P_\rho^R| = \sqrt{|g_1|^2 |g_2|^2}. \quad (47)$$

Evidently, the pixel magnitude of the near-zone camera is the geometric mean of the two traditional beamformer output powers. $|P_\rho^R|$ has better resolution than $|g_1|^2$ and better anisotropic sensitivity than $|g_2|^2$.

Image formation with alternative light fields uses the conjugate field and field measurements to aggregate and isolate power in different ways. In general, image pixel values do not neatly factor into the product of beamformers, as they do with the Rihaczek.

6 Concluding Remarks

We enable the use of existing incoherent imaging tools for coherent imaging applications, by extending the light field to coherent radiation. We explain how to formulate, capture, and form images from quasi light fields. By synthesizing existing research in optics, quantum physics, and signal processing, we motivate quasi light fields, show how quasi light fields extend the traditional light field, and characterize the properties of different quasi light fields. We explain why capturing quasi light fields directly with intensity measurements is inherently limiting, and demonstrate via simulation how processing scalar field measurements in different ways leads to a rich set of energy localization tradeoffs. We show how coherent image formation using quasi light fields is complicated by an implicit far-zone (Fraunhofer) assumption and the fact that not all quasi light fields are constant along rays. We demonstrate via simulation that a pure light field representation is incapable of modeling near-zone diffraction effects, but that quasi light fields can be augmented with a distance parameter for greater near-zone imaging accuracy. We show how image formation using light fields generalizes the classic beamforming algorithm, allowing for new tradeoffs between resolution and anisotropic sensitivity.

Although we have assumed perfectly coherent radiation, tools from partial coherence theory 1) allow us to generalize our results, and 2) provide an alternative perspective on image formation. First, our results extend to broadband radiation of any state of partial coherence by replacing $U(\mathbf{r}^R)U^*(\mathbf{r}^C)$ with the cross-spectral density $W(\mathbf{r}^R, \mathbf{r}^C, \nu)$. W provides a statistical description of the radiation, indicating how light at two different positions, \mathbf{r}^R and \mathbf{r}^C , is correlated at each frequency ν [38]. Second, W itself may be propagated along rays in an approximate asymptotic sense [39, 40], which forms the basis of an entirely different framework for using rays for image formation, using the cross-spectral density instead of the light field as the core representation.

We present a model of coherent image formation that strikes a balance between utility and comprehensive predictive power. On the one hand, quasi light fields offer more options and tradeoffs than their traditional, incoherent counterpart. In this manner, the connection

between quasi light fields and quasi-probability distributions in quantum physics reminds us of the potential benefits of forgoing a single familiar tool in favor of a multitude of useful yet less familiar ones. On the other hand, compared with Maxwell’s equations, quasi light fields are less versatile. Therefore, quasi light fields are attractive to researchers who desire more versatility than traditional energy-based methods, yet a more specialized model of image formation than Maxwell’s equations.

Quasi light fields illustrate the limitations of the simple definition of image formation ubiquitous in incoherent imaging. An image is the visualization of some underlying physical reality, and the energy emitted from a portion of a scene surface towards a virtual aperture is not a physically precise quantity when the radiation is coherent, according to classical electromagnetic wave theory. Perhaps a different image definition may prove more fundamental for coherent imaging, or perhaps a quantum optics viewpoint is required for precision. Although we have borrowed the mathematics from quantum physics, our entire discussion has been classical. Yet if we introduce quantum optics and the particle nature of light, we may unambiguously speak of the probability that a photon emitted from a portion of a scene surface is intercepted by a virtual aperture.

Acknowledgements

This work was supported, in part, by Microsoft Research, MIT Lincoln Laboratory, and Semiconductor Research Corporation through the FCRP Center for Circuit & System Solutions (C2S2).

References

- [1] M. Levoy and P. Hanrahan, “Light field rendering,” in *Proceedings of ACM SIGGRAPH 96*, (ACM, 1996), pp. 31–42.
- [2] R. Ng, M. Levoy, M. Brédif, G. Duval, M. Horowitz, and P. Hanrahan, “Light field photography with a hand-held plenoptic camera,” Tech. Rep. CTSR 2005-02, Stanford University, CA (2005).
- [3] W. Chun and O. S. Cossairt, “Data processing for three-dimensional displays,” United States Patent 7,525,541 (2009).
- [4] R. Ziegler, S. Bucheli, L. Ahrenberg, M. Magnor, and M. Gross, “A bidirectional light field – hologram transform,” *Computer Graphics Forum* **26**, 435–446 (2007).
- [5] T. L. Szabo, *Diagnostic Ultrasound Imaging: Inside Out* (Elsevier Academic Press, 2004).
- [6] M. J. Bastiaans, “Application of the Wigner distribution function in optics,” in *The Wigner Distribution - Theory and Applications in Signal Processing*, W. Mecklenbräuker and F. Hlawatsch, eds. (Elsevier Science B.V., 1997), pp. 375–426.

- [7] A. Walther, “Radiometry and coherence,” J. Opt. Soc. Am. **58**, 1256–1259 (1968).
- [8] E. Wigner, “On the quantum correction for thermodynamic equilibrium,” Phys. Rev. **40**, 749–759 (1932).
- [9] J. Ville, “Théorie et applications de la notion de signal analytique,” Cables et Transmission **2**, 61–74 (1948).
- [10] K. D. Stephan, “Radiometry before World War II: Measuring infrared and millimeter-wave radiation 1800-1925,” IEEE Antennas and Propagation Magazine **47**, 28–37 (2005).
- [11] S. Chandrasekhar, *Radiative Transfer* (Dover Publications, 1960).
- [12] A. T. Friberg, G. S. Agarwal, J. T. Foley, and E. Wolf, “Statistical wave-theoretical derivation of the free-space transport equation of radiometry,” J. Opt. Soc. Am. B **9**, 1386–1393 (1992).
- [13] A. Gershun, “The light field,” J. Math. Phys. **18**, 51–151 (1939).
- [14] M. Born and E. Wolf, *Principles of Optics*, 7th ed. (Cambridge University Press, 1999).
- [15] A. Walther, “Radiometry and coherence,” J. Opt. Soc. Am. **63**, 1622–1623 (1973).
- [16] E. Wolf, “Coherence and radiometry,” J. Opt. Soc. Am. **68**, 6–17 (1978).
- [17] G. S. Agarwal, J. T. Foley, and E. Wolf, “The radiance and phase-space representations of the cross-spectral density operator,” Opt. Commun. **62**, 67–72 (1987).
- [18] E. H. Adelson and J. R. Bergen, “The plenoptic function and the elements of early vision,” in *Computational Models of Visual Processing*, M. S. Landy and J. A. Movshon, eds. (MIT Press, 1991), pp. 3–20.
- [19] S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. F. Cohen, “The lumigraph,” in *Proceedings of ACM SIGGRAPH 96*, (ACM, 1996), pp. 43–54.
- [20] B. Boashash, ed., *Time Frequency Signal Analysis and Processing* (Elsevier, 2003).
- [21] R. W. Boyd, *Radiometry and the Detection of Optical Radiation* (John Wiley & Sons, 1983).
- [22] A. Adams and M. Levoy, “General linear cameras with finite aperture,” in *Proc. Eurographics Symposium on Rendering*, (Eurographics, 2007).
- [23] J. W. Goodman, *Introduction to Fourier Optics* (McGraw-Hill, 1968).
- [24] A. T. Friberg, “On the existence of a radiance function for finite planar sources of arbitrary states of coherence,” J. Opt. Soc. Am. **69**, 192–198 (1979).

- [25] P. Flandrin, *Time-Frequency / Time-Scale Analysis* (Academic Press, 1999).
- [26] D. J. Griffiths, *Introduction to Quantum Mechanics* (Pearson Education, 2005).
- [27] G. S. Agarwal and E. Wolf, “Calculus for functions of noncommuting operators and general phase-space methods in quantum mechanics. I. Mapping theorems and ordering of functions of noncommuting operators,” *Phys. Rev. D* **2**, 2161–2186 (1970).
- [28] J. T. Foley and E. Wolf, “Radiometry as a short-wavelength limit of statistical wave theory with globally incoherent sources,” *Opt. Commun.* **55**, 236–241 (1985).
- [29] J. R. Guerci, “Theory and application of covariance matrix tapers for robust adaptive beamforming,” *IEEE Trans. Signal Process.* **47**, 977–985 (1999).
- [30] Z. Zhang and M. Levoy, “Wigner distributions and how they relate to the light field,” in *Proceedings of ICCP 09*, (IEEE, 2009).
- [31] J. G. Kirkwood, “Quantum statistics of almost classical assemblies,” *Phys. Rev.* **44**, 31–37 (1933).
- [32] A. Rihaczek, “Signal energy distribution in time and frequency,” *IEEE Trans. Information Theory* **14**, 369–374 (1968).
- [33] ZEMAX Development Corporation, Bellevue, WA, *Optical Design Program User’s Guide* (2006).
- [34] M. A. Alonso, “Radiometry and wide-angle wave fields. I. Coherent fields in two dimensions,” *J. Opt. Soc. Am. A* **18**, 902–909 (2001).
- [35] R. G. Littlejohn and R. Winston, “Corrections to classical radiometry,” *J. Opt. Soc. Am. A* **10**, 2024–2037 (1993).
- [36] J. R. Shewell and E. Wolf, “Inverse diffraction and a new reciprocity theorem,” *J. Opt. Soc. Am.* **58**, 1596–1603 (1968).
- [37] H. L. Van Trees, *Optimum Array Processing* (John Wiley & Sons, 2002).
- [38] L. Mandel and E. Wolf, *Optical Coherence and Quantum Optics* (Cambridge University Press, 1995).
- [39] A. M. Zysk, P. S. Carney, and J. C. Schotland, “Eikonal method for calculation of coherence functions,” *Phys. Rev. Lett.* **95**, 043904 (2005).
- [40] R. W. Schoonover, A. M. Zysk, P. S. Carney, J. C. Schotland, and E. Wolf, “Geometrical optics limit of stochastic electromagnetic fields,” *Phys. Rev. A* **77**, 043831 (2008).